

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 August 2001 (30.08.2001)

PCT

(10) International Publication Number
WO 01/63856 A1

- (51) International Patent Classification⁷: **H04L 12/56, H04Q 7/22**
- (21) International Application Number: **PCT/SE01/00407**
- (22) International Filing Date: **23 February 2001 (23.02.2001)**
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
60/184,975 25 February 2000 (25.02.2000) US
09/698,786 27 October 2000 (27.10.2000) US
- (71) Applicant: **TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE).**
- (72) Inventors: **SCHULTZ, Göran; Björkhagsgatan 10, FIN-21600 Pargas (FI). PEISA, Janne; Metsäpirtintie 12D17, FIN-02130 Espoo (FI). WIGELL, Toomas; Venepuuntie 6, FIN-02360 Espoo (FI). MATINMIKKO, Reijo; Veininkatu 41 A 4, FIN-02730 Espoo (FI).**
- (74) Agent: **MAGNUSSON, Monica; Ericsson Radio Systems AB, Patent Unit Radio Access, S-164 80 Stockholm (SE).**
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— *with international search report*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: **FLOW CONTROL BETWEEN TRANSMITTER AND RECEIVER ENTITIES IN A COMMUNICATIONS SYSTEM**

(57) Abstract: In certain wireless communication systems, a common channel MAC-c layer (500), a dedicated channel MAC-d layer (510), and a Radio Link Controller are respectively located in a Radio Network Control. The MAC-c layer (500) is provided with a flow control mechanism (530) for managing a number of MAC-d entity data traffic flows directed to the MAC-c buffer (525), so that buffer fill level and the MAC-c entity (500) will remain at a desired level, and all respective data flows each gets a fair share of the data rate between the MAC-c (500) and MAC-d (510) entities. The flow control operates to share MAC-c buffer space with each MAC-d entity (510) providing an active data flow in a sequential, round robin fashion or based on flow activity (greedy manner). Buffer space is assigned or allocated on the basis of MAC-c upper fill level, and in some embodiments also on the basis of active MAC-d buffer fill levels.

WO 01/63856 A1

FLOW CONTROL BETWEEN TRANSMITTER AND RECEIVER
ENTITIES
IN A COMMUNICATIONS SYSTEM

5 **CROSS-REFERENCES TO RELATED APPLICATIONS**

 This Non-provisional Application for Patent claims the benefit of priority from, and hereby incorporates by reference the entire disclosure of, co-pending U.S. Provisional Application for Patent Serial No. 60/184,975, (Attorney Docket No. 34646-00458USPL) filed February 25, 2000. Co-pending U.S. Provisional
10 Applications for Patent Serial Nos. 60/185,005 (Attorney Docket No. 34646-00459USPL) and 60/185,003 (Attorney Docket No. 34646-00460USPL), both filed on February 25, 2000, are also hereby incorporated by reference in their entirety herein.

 This Non-provisional Application for Patent is related by subject matter to U.S.
15 Non-provisional Applications for Patent Nos. 09/698785 (Attorney Docket No. 34646-00459USPT) and 09/698672 (Attorney Docket No. 34646-00460USPT), both of which are filed on even date herewith. These two U.S. Non-provisional Applications for Patent are also hereby incorporated by reference in their entirety herein.

20 **BACKGROUND OF THE INVENTION**

Technical Field of the Invention

 The present invention relates in general to the field of communications systems, and in particular, by way of example but not limitation, to handling potential buffer overflows resulting from users overloading a communications system.

25 **Description of Related Art**

 Access to and use of wireless networks is becoming increasingly important and popular for business, social, and recreational purposes. Users of wireless networks now rely on them for both voice and data communications. Furthermore, an ever increasing number of users demand both an increasing array of services and
30 capabilities as well as greater bandwidth for activities such as Internet surfing. To

-2-

address and meet the demands for new services and greater bandwidth, the wireless communications industry constantly strives to improve the number of services and the throughput of their wireless networks. Expanding and improving the infrastructure necessary to provide additional services and higher bandwidth is an expensive and manpower-intensive undertaking. Moreover, high-bandwidth data streams will eventually be demanded by consumers to support features such as real-time audio-visual downloads and live audio-visual communication between two or more people. In the future, it will therefore become necessary and/or more cost-effective to introduce next generation wireless system(s) instead of attempting to upgrade existing system(s).

To that end, the wireless communications industry intends to continue to improve the capabilities of the technology upon which it relies and that it makes available to its customers by deploying next generation system(s). Protocols for a next-generation standard that is designed to meet the developing needs of wireless customers is being standardized by the 3rd Generation Partnership Project (3GPP). The set of protocols is known collectively as the Universal Mobile Telecommunications System (UMTS).

Referring now to FIG. 1, an exemplary wireless communications system with which the present invention may be advantageously employed is illustrated generally at 100. In a UMTS network 100, the network 100 includes a core network 120 and a UMTS Terrestrial Radio Access Network (UTRAN) 130. The UTRAN 130 is composed of, at least partially, a number of Radio Network Controllers (RNCs) 140, each of which may be coupled to one or more neighboring Node Bs 150. Each Node B 150 is responsible for a given geographical cell and the controlling RNC 140 is responsible for routing user and signaling data between that Node B 150 and the core network 120. All of the RNCs 140 may be directly or indirectly coupled to one another. A general outline of the UTRAN 130 is given in Technical Specification TS 25.401 V2.0.0 (1999-09) of the 3rd Generation Partnership Project, 3GPP, which is hereby incorporated by reference in its entirety herein. The UMTS network 100 also includes multiple user equipments (UEs) 110. UE may include, for example, mobile

-3-

stations, mobile terminals, laptops/personal digital assistants (PDAs) with wireless links, etc.

In an exemplary second layer of the architectural structure of the UMTS, a set of Radio Access Bearers (RABs) is provided to make radio resources and services available to user applications. For each mobile station, there may be one or several RABs, and data flows from the RABs to respective Radio Link Control (RLC) entities in the form of segments. The RLC entities buffer the received data segments and map the RABs onto respective logical channels. A Medium Access Control (MAC) entity receives data transmitted in the logical channels, and further maps the data from the logical channels onto a set of transport channels. The transport channels are, in turn, mapped to a single physical transport channel having a specified total band width allocated to it by the associated network. A MAC entity connected to a dedicated transport channel is known as MAC-d, and a MAC entity connected to a common transport channel is known as MAC-c. Preferably, there is one MAC-d entity for each mobile station in the UMTS, and one MAC-c entity for each cell.

During each transmission time interval of the system, the MAC entity has to decide how much data to transmit on each transport channel connected to it. In making this decision, it is necessary to share the total available band width between the logical channels receiving information from different RABs and their respective RLCs coupled to different mobile stations. In the past, sharing resources between multiple input data flows has been referred to as Generalized Processor Sharing (GPS). However, it has been recognized that there are difficulties involved in applying GPS directly to bandwidth allocations in a UMTS network. More specifically, GPS assumes that data can be sent on the MAC entity logical channels in infinitely small blocks. This is not possible in UMTS. Accordingly, it is necessary to provide an alternative arrangement for managing or controlling data flow in the UMTS, in order to share the total available bandwidth among different logical channels receiving information from different RAB's.

30

SUMMARY OF THE INVENTION

The above-identified deficiencies, as well as others that are associated with existing schemes, are remedied by the methods, systems and arrangements of the present invention. More particularly, an arrangement is made available which provides data
5 flow from a number of different MAC-d entities, respectfully associated with different logical channels, with a fair share of the data rate between the Mac-d entities and a common MAC-c entity, and which also maintains the MAC-c buffer at or close to a desired fill level. This is achieved by sharing MAC-c buffer free space between active data flows of Mac-d entities according to a sequential, round robin or flow activity
10 principle and using available information pertaining to the buffer fill levels of the respective RLC/MAC-d entities.

Generally, the invention is directed to a method for providing flow control in a wireless communications system which comprises the steps of providing a receiver entity and one or more transmitter entities, a given transmitter entity transmitting when
15 it receives credit from the receiver entity. The method further comprises the step of determining whether to give credit to respective transmitter entities.

In a preferred embodiment of the invention, the method comprises transmitting packets of data from respective transmitter entities to the receiver entity in a sequential or round robin fashion. The receiver entity comprises a common Medium Access
20 Control Entity (MAC-c) having a receiver buffer, and each of the transmitter entities comprises a dedicated Medium Access Control (Mac-d) entity, each transmitter entity bearing associated with another buffer. In one useful embodiment, the determination of the determining step is based, at least in part, on the fill level of the receiver buffer, and also on the fill levels of one or more buffers associated with respective transmitter
25 entities. In another useful embodiment, determination is based, at least in part, on the fill level of the receiver buffer, but not on the fill levels of any of the buffers associated with transmitter entities. Another embodiment of the invention includes the steps of ascertaining if a first active transmitter entity has a current credit which is less than the buffer fill level of an associated buffer, and if so, increasing the current credit of the
30 first active transmitter entity by one and decreasing the available credit amount by one. Such ascertaining step is repeated for each active transmitter entity until either the

-5-

available credit amount is exhausted, or all active transmitter entities have their corresponding current credits to their corresponding buffer fill levels. Such embodiment further includes the step of recording the last active transmitter entity to receive a credit, and making another active transmitter entity immediately following the last active transmitter in a sequence to be the first active transmitter entity to receive a credit on the next occasion when credits are distributed.

The above-described and other features of the present invention are explained in detail hereinafter with reference to the illustrative examples shown in the accompanying drawings. Those skilled in the art will appreciate that the described embodiments are provided for purposes of illustration and understanding and that numerous equivalent embodiments are contemplated herein.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the methods, systems, and arrangements of the present invention may be had by reference to the following detailed description when taken in conjunction with the accompanying drawings wherein:

FIG. 1 illustrates an exemplary wireless communications system with which the present invention may be advantageously employed;

FIG. 2 illustrates a protocol model for an exemplary next-generation system with which the present invention may be advantageously employed;

FIG. 3 illustrates a view of an exemplary second layer architecture of an exemplary next-generation system in accordance with the present invention;

FIG. 4 illustrates an exemplary method in flowchart form for allocating bandwidth resources to data flow streams between entities in the exemplary second layer architecture of FIG. 3;

FIG. 5 illustrates a flow control comprising an embodiment of the invention interconnected between transmit and receive entities.

FIG. 6 shows a portion of the embodiment of Figure 5 in greater detail.

FIG. 7 illustrates a buffer for the embodiment of Figure 5.

FIG. 8 illustrates operation of the flow control of the embodiment shown in Figure 5.

-6-

FIG 9 shows a buffer for a second embodiment of the invention.

FIG. 10 is a schematic diagram illustrating operation of the second embodiment.

FIG. 11 is a graph depicting the relationship between buffer fill level and credits for the second embodiment.

FIGs. 12 and 13 are schematic diagrams for illustrating a third embodiment of the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, for purposes of explanation and not limitation, specific details are set forth, such as particular circuits, logic modules (implemented in, for example, software, hardware, firmware, some combination thereof, etc.), techniques, etc. in order to provide a thorough understanding of the invention. However, it will be apparent to one of ordinary skill in the art that the present invention may be practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well-known methods, devices, logical code (e.g., hardware, software, firmware, etc.), etc. are omitted so as not to obscure the description of the present invention with unnecessary detail.

A preferred embodiment of the present invention and its advantages are best understood by referring to FIGS. 1-13 of the drawings, like numerals being used for like and corresponding parts of the various drawings. Aspects of the UMTS are used to describe a preferred embodiment of the present invention. However, it should be understood that the principles of the present invention are applicable to other wireless communication standards (or systems), especially those in which communication is packet-based.

Referring now to FIG. 2, a protocol model for an exemplary next-generation system with which the present invention may be advantageously employed is illustrated generally at 200. In the protocol model 200 (e.g., for a Forward Access Channel (FACH) transport Channel type), the "Uu" indicates the interface between UTRAN 130 and the UE 110, and "Iub" indicates the interface between the RNC 140 and a Node B 150 (where "Node B" is a generalization of, for example, a Base

-7-

Transceiver Station (BTS)). User and signaling data may be carried between an RNC 140 and a UE 110 using Radio Access Bearers (RABs) (as illustrated hereinbelow with reference to FIG. 3). Typically, a UE 110 is allocated one or more RABs, each of which is capable of carrying a flow of user or signaling data. RABs are mapped onto
5 respective logical channels. At the Media Access Control (MAC) layer, a set of logical channels is mapped in turn onto a transport Channel, of which there are two types: a "common" transport Channel which is shared by different UEs 110 and a "dedicated" transport Channel which is allocated to a single UE 110 (thus leading to the terms "MAC-c" and "MAC-d"). One type of common Channel is the FACH. A
10 basic characteristic of a FACH is that it is possible to send one or more fixed size packets per transmission time interval (e.g., 10, 20, 40, or 80 ms). Several transport channels (e.g., FACHs) are in turn mapped at the physical layer onto a Secondary Common Control Physical CHANNEL (S-CCPCH) for transmission over the air interface between a Node B 150 and a UE 110.

15 When a UE 110 registers with an RNC 140 via a Node B 150, that RNC 140 acts at least initially as both the serving and the controlling RNC 140 for the UE 110. (The serving RNC 140 may subsequently differ from the controlling RNC 140 in a UMTS network 100, but the presence or absence of this condition is not particularly relevant here.) The RNC 140 both controls the air interface radio resources and
20 terminates the layer 3 intelligence (e.g., the Radio Resource Control (RRC) protocol), thus routing data associated with the UE 110 directly to and from the core network 120.

It should be understood that the MAC-c entity in the RNC 140 transfers MAC-c Packet Data Units (PDUs) to the peer MAC-c entity at the UE 110 using the services
25 of the FACH Frame Protocol (FACH FP) entity between the RNC 140 and the Node B 150. The FACH FP entity adds header information to the MAC-c PDUs to form FACH FP PDUs which are transported to the Node B 150 over an AAL2 (or other transport mechanism) connection. An interworking function at the Node B 150 interworks the FACH frame received by the FACH FP entity into the PHY entity.

30 In an exemplary aspect of the scenario illustrated in FIG. 2, an important task of the MAC-c entity is the scheduling of packets (MAC PDUs) for transmission over

-8-

the air interface. If it were the case that all packets received by the MAC-c entity were of equal priority (and of the same size), then scheduling would be a simple matter of queuing the received packets and sending them on a first come first served basis (e.g., first-in, first-out (FIFO)). However, UMTS defines a framework in which different Quality of Services (QoS) may be assigned to different RABs. Packets corresponding to a RAB that has been allocated a high QoS should be transmitted over the air interface at a high priority whilst packets corresponding to a RAB that has been allocated a low QoS should be transmitted over the air interface at a lower priority. Priorities may be determined at the MAC entity (e.g., MAC-c or MAC-d) on the basis of RAB parameters.

UMTS deals with the question of priority by providing at the controlling RNC 140 a set of queues for each FACH. The queues may be associated with respective priority levels. An algorithm, which is defined for selecting packets from the queues in such a way that packets in the higher priority queues are (on average) dealt with more quickly than packets in the lower priority queues, is implemented. The nature of this algorithm is complicated by the fact that the FACHs that are sent on the same physical Channel are not independent of one another. More particularly, a set of Transport Format Combinations (TFCs) is defined for each S-CCPCH, where each TFC includes a transmission time interval, a packet size, and a total transmission size (indicating the number of packets in the transmission) for each FACH. The algorithm should select for the FACHs a TFC which matches one of those present in the TFC set in accordance with UMTS protocols.

Preferably, a packet received at the controlling RNC 140 is placed in a queue (for transmission on a FACH), where the queue corresponds to the priority level attached to the packet as well as to the size of the packet. The FACH is mapped onto a S-CCPCH at a Node B 150 or other corresponding node of the UTRAN 130. In an alternative preference, the packets for transmission on the FACH are associated with either a Dedicated Control CHANNEL (DCCH) or to a Dedicated Traffic CHANNEL (DTCH). It should be noted that, preferably, each FACH is arranged to carry only one size of packets. However, this is not necessary, and it may be that the packet size that can be carried by a given FACH varies from one transmission time interval to another.

-9-

As alluded to hereinabove, the UE 110 may communicate with the core network 120 of the UMTS system 100 via separate serving and controlling (or drift) RNCs 140 within the UTRAN 130 (e.g., when the UE 110 moves from an area covered by the original serving RNC 140 into a new area covered by a controlling/drift RNC 140) (not specifically shown). Signaling and user data packets destined for the UE 110 are received at the MAC-d entity of the serving RNC 140 from the core network 120 and are "mapped" onto logical channels, namely a Dedicated Control Channel (DCCH) and a Dedicated traffic CHANNEL (DTCH), for example. The MAC-d entity constructs MAC Service Data Units (SDUs), which include a payload section containing logical channel data and a MAC header containing, *inter alia*, a logical channel identifier. The MAC-d entity passes the MAC SDUs to the FACH FP entity. This FACH FP entity adds a further FACH FP header to each MAC SDU, where the FACH FP header includes a priority level that has been allocated to the MAC SDU by an RRC entity. The RRC is notified of available priority levels, together with an identification of one or more accepted packet sizes for each priority level, following the entry of a UE 110 into the coverage area of the drift RNC 140.

The FACH FP packets are sent to a peer FACH FP entity at the drift RNC 140 over an AAL2 (or other) connection. The peer FACH FP entity decapsulates the MAC-d SDU and identifies the priority contained in the FRAME FP header. The SDU and associated priority are passed to the MAC-c entity at the controlling RNC 140. The MAC-c layer is responsible for scheduling SDUs for transmission on the FACHs. More particularly, each SDU is placed in a queue corresponding to its priority and size. For example, if there are 16 priority levels, there will be 16 queue sets for each FACH, with the number of queues in each of the 16 queue sets depending upon the number of packet sizes accepted for the associated priority. As described hereinabove, SDUs are selected from the queues for a given FACH in accordance with some predefined algorithm (e.g., so as to satisfy the TFC requirements of the physical channel).

The scheme described hereinbelow with reference to FIGS. 3 and 4 relates to data transmission in a telecommunications network and in particular, though not necessarily, to data transmission in a UMTS.

-10-

As noted hereinabove, the 3GPP is currently in the process of standardizing a new set of protocols for mobile telecommunications systems. The set of protocols is known collectively as the UMTS. With reference to FIG. 3, a view of an exemplary second layer architecture of an exemplary next-generation system in accordance with the present invention is illustrated generally at 300. Specifically, by way of example only, the exemplary second layer architecture 300 illustrates a simplified UMTS layer 2 protocol structure which is involved in the communication between mobile stations (e.g. mobile telephones), or more broadly UEs 110, and Radio Network Controllers (RNCs) 140 of a UMTS network 100. The RNCs 140 are analogous to the Base Station Controllers (BSCs) of existing GSM mobile telecommunications networks, communicating with the mobile stations via Node Bs 150.

The layer 2 structure of the exemplary second layer architecture 300 includes a set of Radio Access Bearers (RABs) 305 that make available radio resources (and services) to user applications. For each mobile station there may be one or several RABs 305. Data flows (e.g., in the form of segments) from the RABs 305 are passed to respective Radio Link Control (RLC) entities 310, which amongst other tasks buffer the received data segments. There is one RLC entity 310 for each RAB 305. In the RLC layer, RABs 305 are mapped onto respective logical channels 315. A Medium Access Control (MAC) entity 320 receives data transmitted in the logical channels 315 and further maps the data from the logical channels 315 onto a set of transport channels 325. The transport channels 325 are finally mapped to a single physical transport channel 330, which has a total bandwidth (e.g., of <2Mbits/sec) allocated to it by the network. Depending whether a physical channel is used exclusively by one mobile station or is shared between many mobile stations, it is referred to as either a "dedicated physical channel" or a "common channel". A MAC entity connected to a dedicated physical channel is known as MAC-d; there is preferably one MAC-d entity for each mobile station. A MAC entity connected to a common channel is known as MAC-c; there is preferably one MAC-c entity for each cell.

The bandwidth of a transport channel 325 is not directly restricted by the capabilities of the physical layer 330, but is rather configured by a Radio Resource Controller (RRC) entity 335 using Transport Formats (TFs). For each transport

-11-

channel 325, the RRC entity 335 defines one or several Transport Block (TB) sizes. Each Transport Block size directly corresponds to an allowed MAC Protocol Data Unit (PDU) and tells the MAC entity what packet sizes it can use to transmit data to the physical layer. In addition to block size, the RRC entity 335 informs the MAC entity 320 of a Transport Block Set (TBS) size, which is the total number of bits the MAC entity can transmit to the physical layer in a single transmission time interval (TTI). The TB size and TBS size, together with some additional information relating to the allowed physical layer configuration, form a TF. An example of a TF is (TB=80 bits, TBS=160 bits), which means that the MAC entity 320 can transmit two 80 bit packets in a single TTI. Thus, this TF can be written as TF=(80,160). The RRC entity 335 also informs the MAC entity of all possible TFs for a given transport channel. This combination of TFs is called a Transport Format Combination (TFC). An example of a TFC is {TF1=(80, 80), TF2=(80, 160)}. In this example, the MAC entity can choose to transmit one or two PDUs in one TTI on the particular transport channel in question; in both cases, the PDUs have a size of 80 bits.

In each TTI, the MAC entity 320 has to decide how much data to transmit on each transport channel 325 connected to it. These transport channels 325 are not independent of one another, and are later multiplexed onto a single physical channel 330 at the physical layer 330 (as discussed hereinabove). The RRC entity 335 has to ensure that the total transmission capability on all transport channels 325 does not exceed the transmission capability of the underlying physical channel 330. This is accomplished by giving the MAC entity 320 a Transport Format Combination Set (TFCS), which contains the allowed Transport Format Combinations for all transport channels.

By way of example, consider a MAC entity 320 which has two transport channels 325 that are further multiplexed onto a single physical channel 330, which has a transport capacity of 160 bits per transmission time interval (It should be understood that, in practice, the capacity will be much greater than 160). The RRC entity 335 could decide to assign three transport formats TF1=(80, 0), TF2=(80, 80) and TF3=(80, 160) to both transport channels 325. Clearly however, the MAC entity 320 cannot choose to transmit on both transport channels 325 at the same time using

-12-

TF3 as this would result in the need to transmit 320 bits on the physical channel 330, which has only a capability to transmit 160 bits. The RRC entity 335 has to restrict the total transmission rate by not allowing all combinations of the TFs. An example of this would be a TFCS as follows [{{(80, 0), (80, 0)}, {(80, 0), (80, 80)}, {(80, 0), (80, 160)}, {(80, 80), (80, 0)}, {(80, 80), (80, 80)}, {(80, 160), (80, 0)}}], where the transport format of transport channel "1" is given as the first element of each element pair, and the transport format of transport channel "2" is given as the second element. As the MAC entity 320 can only choose one of these allowed transport format combinations from the transport format combination set, it is not possible to exceed the capability of the physical channel 330.

An element of the TFCS is pointed out by a Transport Format Combination Indicator (TFCI), which is the index of the corresponding TFC. For example, in the previous example there are 6 different TFCs, meaning that the TFCI can take any value between 1 and 6. The TFCI=2 would correspond to the second TFC, which is {(80, 0), (80, 80)}, meaning that nothing is transmitted from the first transport channel and a single packet of 80 bits is transmitted from the second transport channel.

It is of course necessary to share the total available bandwidth between the logical channels 315. The decision to distribute the bandwidth to different transport channels is made by the MAC entity 320 for each transmission time interval by choosing a suitable TFCI. This sharing of bandwidth can be done in several ways, for example by giving an absolute preference to flows which are considered to be more important than others. This would be the easiest method to implement, but can result in a very unfair distribution of the bandwidth. Specifically, it is possible that flows that have lower priorities are not allowed to transmit for prolonged periods of time. This can result in extremely poor performance if the flow control mechanism of a lower priority flow reacts to this. A typical example of such a flow control mechanism can be found in the present day Transmission Control Protocol (TCP) protocol used in the Internet. In existing technologies, such as Internet Protocol (IP) and Asynchronous Transfer Mode (ATM) networks, provision is made for allocating resources on a single output channel to multiple input flows. However, the algorithms

-13-

used to share out the resources in such systems are not directly applicable to UMTS where multiple input flows are transmitted on respective logical output channels.

Sharing resources between multiple input data flows is referred to as Generalized Processor Sharing (GPS). This GPS, when employed in systems having only a single output channel, is known as Weighted Fair Queuing (WFQ) and is described in a paper entitled "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case", A. K. Parekh, R. G. Gallager, published in IEEE/ACM Transactions On Networking, Vol. 1, No. 3, June 1993, Pp. 344-357. Stated simply, GPS involves calculating a GPS weight for each input flow on the basis of certain parameters associated with the flow. The weights calculated for all of the input flows are added together, and the total available output bandwidth is divided amongst the input flows depending upon the weight of each flow as a fraction of the total weight. GPS could be applied to the MAC entity in UMTS, with the weighting for each input flow being determined (by the RRC entity) on the basis of certain RAB parameters, which are allocated to the corresponding RAB by the network. In particular, an RAB parameter may equate to a Quality of Service (QoS) or Guaranteed rate allocated to a user for a particular network service.

Continuing now with the scheme described herein with reference to FIGS. 3 and 4, it was realized that there are difficulties involved in applying GPS directly to bandwidth allocation in a UMTS network, as GPS assumes that data can be sent on the MAC entity logical channels in infinitely small blocks. This is not possible in UMTS, as UMTS relies upon Transport Format Combinations Sets (TFCSSs) as the basic mechanism defining how much data can be sent in each TTI. If GPS is to be employed in UMTS, it is necessary to select the TFC (from the TFCS) which most closely matches the bandwidth allocated to an input flow by GPS. The result of this approach is that the actual amount of data sent for an input stream in a given frame either may fall below the optimized rate or may exceed that optimized rate. In the former case, a backlog of unsent data may build up for the input flow.

It is an object of the scheme described herein with reference to FIGS. 3 and 4 to overcome, or at least mitigate, the disadvantage noted in the preceding paragraph. This and other objects are achieved at least in part by maintaining a backlog counter

-14-

which keeps track of the backlog of unsent data for a given input flow to the MAC entity. The backlog is taken into account when determining an appropriate TFC for that input flow for a subsequent frame. According to a first aspect of this scheme, there is provided a method of allocating transmission resources at a Media Access Control (MAC) entity of a node of a Universal Mobile Telecommunications System (UMTS), the method including the following steps for each frame of an output data flow: computing for each input flow to the MAC entity a fair share of the available output bandwidth of the MAC entity; selecting a Transport Format Combination (TFC) from a TFC Set (TFCS) on the basis of the bandwidth share computed for the input flows, where the TFC includes a Transport Format allocated to each input flow; and for each input flow, if the allocated TF results in a data transmission rate which is less than the determined fair distribution, adding the difference to a backlog counter for the input flow, where the value of the backlog counter(s) is taken into account when selecting a TFC for the subsequent frame of the output data flow. Embodiments of this scheme allow the TFC selection process for a subsequent frame to take into account any backlogs which exists for the input flows. The tendency is to adjust the selected TFC to reduce the backlogs. Such a backlog may exist due to the finite number of data transmission possibilities provided for by the TFCS. Nodes at which the method of this scheme may be employed include mobile stations (such as mobile telephones and communicator type devices) (or more generally UEs) and Radio Network Controllers (RNCs).

Preferably, the input flows to the MAC entity are provided by respective Radio Link Control (RLC) entities. Also preferably, each RLC entity provides buffering for the associated data flow. Also preferably, the step of computing a fair share of resources for an input flow is carried out by a Radio Resource Control (RRC) entity. Also preferably, the step of computing a fair share of resources for an input flow includes the step of determining the weighting given to that flow as a fraction of the sum of the weights given to all of the input flows. The fair share may then be determined by multiplying the total output bandwidth by the determined fraction. Also preferably, this step may involve using the Generalized Processor Sharing (GPS) mechanism. The weighting for a data flow may be defined by one or more Radio

-15-

Access Bearer (RAB) parameters allocated to a RAB by the UMTS network, where the RAB is associated with each MAC input flow. Also preferably, in the event that the backlog counter for a given input flow has a positive value, the method further includes the step of adding the value of the backlog counter to the computed fair share for that flow and selecting a TFC on the basis of the resulting sums for all of the input flows.

In certain embodiments of the scheme described herein with reference to FIGS. 3 and 4, where, for a given input flow, the allocated TF results in a data transmission rate that is more than the determined fair distribution, the difference may be subtracted from the backlog counter for the input flow. According to a second aspect of this scheme, there is provided a node of a Universal Mobile Telecommunications System (UMTS), the node including: a Media Access Control (MAC) entity for receiving a plurality of input data flows; first processor means for computing for each input flow to the MAC entity a fair share of the available output bandwidth of the MAC entity and for selecting a Transport Format Combination (TFC), from a TFC Set (TFCS), on the basis of the bandwidth share computed for the input flows, where the TFC includes a Transport Format allocated to each input flow; second processor means for adding to a backlog counter associated with each input flow the difference between the data transmission rate for the flow resulting from the selected TFC and the determined fair share, if the data transmission rate is less than the determined fair share, where the first processor means is arranged to take into account the value of the backlog counters when selecting a TFC for the subsequent frame of the output data flow. Preferably, the first and second processor means are provided by a Radio Resource Control (RRC) entity.

As is described herein with reference to FIG. 3, a simplified UMTS layer 2 includes one Radio Resource Control (RRC) entity, a Medium Access Control (MAC) entity for each mobile station, and a Radio Link Control (RLC) entity for each Radio Access Bearer (RAB). The MAC entity performs scheduling of outgoing data packets, while the RLC entities provide buffers for respective input flows. The RRC entity sets a limit on the maximum amount of data that can be transmitted from each flow by assigning a set of allowed Transport Format Combinations (TFC) to each MAC

-16-

(referred to as a TFC Set or TFCS), but each MAC must independently decide how much data is transmitted from each flow by choosing the best available Transport Format Combination (TFC) from the TFCS.

With reference now to FIG. 4, an exemplary method in flowchart form for allocating bandwidth resources to data flow streams between entities in the exemplary second layer architecture of FIG. 3 is illustrated generally at 400. The flowchart 400 is a flow diagram of a method of allocating bandwidth resources to, for example, the input flow streams of a MAC entity of the layer 2 of FIG 3. Generally, an exemplary method in accordance with the flowchart 400 may follow the following steps. First, input flows are received at RLCs and the data is buffered (step 405). Information on buffer fill levels is passed to the MAC entity (step 410). After the information on buffer fill levels is passed, the fair MAC bandwidth share for each input flow is computed (step 415). The computed fair share of each is then adjusted by adding the contents of an associated backlog counter to the respective computed fair share (step 420). Once the computed fair shares have been adjusted, a TFC is selected from the TFC set to most closely match the adjusted fair shares (step 425). The RLC is next instructed to deliver packets to the MAC entity according to the selected TFC (step 430). The MAC entity may also schedule packets in accordance with the selected TFC (step 435). After packet scheduling, the traffic channels may be transported on the physical Channel(s) (step 440). Once packet traffic has been transported, the backlog counters should be updated (step 445). The process may continue (via arrow 450) when new input flows are received at the RLCs, which buffer the data (at step 405).

Furthermore, certain embodiment(s) of the scheme operate by calculating at the MAC entity, on a per Transmission Time Interval (TTI) basis, the optimal distribution of available bandwidth using the Generalized Processor Sharing (GPS) approach (*See, e.g., the article by A. K. Parekh et al. referenced hereinabove.*) and by keeping track of how far behind each flow is from the optimal bandwidth allocation using respective backlog counters. The available bandwidth is distributed to flows by using the standard GPS weights, which may be calculated by the RRC using the RAB parameters.

-17-

The method may first calculate the GPS distribution for the input flows and add to the GPS values the current respective backlogs. This is performed once for each 10ms TTI and results in a fair transmission rate for each flow. However, this rate may not be optimal as it may happen that there is not enough data to be sent in all buffers. In order to achieve optimal throughput as well as fairness, the fair GPS distribution is reduced so as to not exceed the current buffer fill level or the maximum allowed rate for any logical Channel. A two step rating process is then carried out.

First, the set of fair rates computed for all of the input flows is compared against possible Transport Format Combinations (TFCs) in turn, with each TFC being scored according to how close it comes to sending out the optimal rate. In practice this is done by simply counting how much of the fair configuration a TFC fails to send (if a given TFC can send all packets at the fair rate, it is given a score of zero) and then considering only the TFCs which have the lowest scores. The closest match is chosen and used to determine the amount of packets sent from each queue. TFCs having an equal score are given a bonus score according to how many extra bits they can send (this can be further weighted by a Quality of Service rating in order to ensure that the excess capacity goes to the bearer with the highest quality class). The final selection is based on a two-level scoring: the TFC with the lowest score is taken. If there are several TFCs with an equal score, the one with the highest bonus score is chosen. This ensures that the rate for each TTI is maximized. Fairness is achieved by checking that if the chosen TFC does not give all flows at least their determined fair rate, the missing bits are added to a backlog counter of the corresponding flow and the selection is repeated for the next TTI. If any of the flows has nothing to transmit, the backlog is set to zero.

This algorithm can be shown to provide bandwidth (and, under certain assumptions, delay bounds) that is close to that of GPS. However, it remains fair and maintains isolation between all flows. It is also computationally simpler than Weighted Fair Queuing algorithms because it utilizes the fact that the MAC layer can transmit on several transport channels at the same time. This results in optimal or close to optimal utilization of the radio interface in the UMTS radio link. The

-18-

following pseudo-code is an outline of an exemplary algorithm for implementing the scheme described hereinabove with reference to FIGS. 3 and 4:

```
/*
5  * GPS based TFC selection. Schedules packets by optimizing the throughput
  * while still keeping the fairness (i.e. guaranteed rates).

int sched_gps() {

10     double weight, weight_sum;
        double score, bonus_score;
        double min_score = HUGE_NUMBER;
        double max_bonus_score = 0;

15     int maxrate;
        int i, j;
        int tfc, tfci, qf, rate, trch;
        int tfc_to_use;

20     double backlog[MAX_TRCH];
        double gps_req[MAX_TRCH];
        double gps_req_comp[MAX_TRCH];

        /* First calculate the sum of the weights of all active queues */

25     weight_sum=0;
        for (trch = 0; trch < MAX_TRCH; trch++) {
            if (queue_fill_state[trch] > 0) {
                weight_sum. += weight_vector[trch];

30         }
    }
}
```

-19-

```

/* Then calculate the fair distribution of available bandwidth
* using GPS. Modify the GPS scheduling reducing the rate if there
* is not enough data in the buffers or if the scheduled rate is
* higher that the maximum rate for a given logical channel
5 */

int gps_rate=0;
for(trch = 0; trch < MAX_TRCH; trch++) {

10     if (queue_fill_state [trch] == 0) {
        backlog[trch] = 0;
    }

    // here we calculate how many bits we should send on each channel
15 // by GPS

    gps_req[trch] = 0;
    gps_req_comp[trch] = 0;
    if(queue_fill_state[trch] > 0) {
20         weight = weight_vector[trch];
        gps_req[trch] = weight / weight_sum * maxrate +
        backlog[trch];
        gps_req_comp[trch] = gps_req[trch];
        if (gps_req_comp[trch] > queue_fill_state[trch]) {
25             gps_req_comp[trch] = queue_fill_state[trch];
        }

        if (gps_req_comp[trch] > trch_max_rate[trch]) {
            gps_req_comp[trch] = trch_max_rate[trch];
        }

30     }
}

```

-20-

```

/* Now we have our basis for selecting the TFC. Score all available
* TFCs by calculating how far they are from the modified GPS
* result. If there are several TFCs that can send the whole GPS
* result (or are
5  * equally close) choose the one that maximises the throughput of
* highest QoS class. Note that the TFCIs are assumed to be in
* increasing order regarding the bandwidth usage
*/

10  for (tfci = 0; tfci < MAX_TPCI; tfci++) {
        rate = score = bonus_score = 0;
        for (trch = 0; trch < MAX_TRCH; trch++) {

                int tbs = tfcs [trch] [tfci] [0];
                int tbss = tfcs [trch] [tfci] [1];
                rate += tbss;

                if (tbss < gps_req_comp[trch]) {
                        score += gps_req_comp [trch] - tbss;
                } else {
                        if (tbss <= queue_fill_state(trch)) {
                                bonus_score += QoS_vector[trch]*(tbss -
20  gps_req_comp[trch]);
                        }
                }

25  }
        }

        if (score < min_score) {
                tfc_to_use = tfci;
                min_score = tfcScore;
                max_bonus_score = bonus_score;
30  }
}

```

-21-

```

    if (score == min_score && bonus_score > max-bonus-score) {
        tfc_to_use = tfci;
        min_score = score;
        max_bonus_score = bonus_score;
5      }
    }

    /* Now we have chosen the TFC to use. Update the backlog and output the
    * right TFCI
    */

10   for (trch = 0; trch < MAX_TRCH; trch++) {

        tbss = tfcs [trch] [tfcToUse] [1];
        if (tbss < queue_fill_state) {
15             if (gps_req [trch] - gps_req_comp[trch]) {
                    backlog [trch] = gpsReq [trch] - tbss;
                    if (backlog[trch] < 0) backlog[trch] = 0;
            } else {
                    backlog[trchG1] = 0;
20             }
        }
    }
    return tfc_to_use;
}

```

25 Referring now to Figure 5, there is shown a portion of the second layer architecture depicted in Figure 3. More particularly, there is shown the MAC entity 320 of Figure 3 comprising a MAC-c entity 500, as described above, connected to a dedicated MAC-d entity 510, as likewise described above. Figure 5 further shows RLC entities 515, which receive data in the form of segments from respective RAB's
30 and map the data onto corresponding logical channels 315. Each of the RLC's 515 is provided with a buffer 520 for buffering received data segments comprising PDU's.

-22-

MAC-d entity 510 serves a single mobile station (not shown) and data transmitted by the MAC-d entity is directed through transport channel 325 to a dedicated physical Channel DPCH.

Referring further to Figure 5, there is shown the common MAC-c entity 500 provided with buffers 525 and connected to receive data, or more specifically PDU's, from MAC-d entity 510 as well as from other MAC-d entities which are not shown. In accordance with the invention, it is desired to provide a flow control mechanism 530 between each of the dedicated MAC-d entities and the common MAC-c entity 500. By selectively managing traffic flow from the MAC-d entities to the buffer 525 of MAC-c entity 500, the flow control 530 enables a number of different MAC-d entities which are respectively in an active data flow mode to share the available free space of MAC-c buffer 525. The flow control is designed to provide each of the active MAC-d entities with a "fair share" of the additional data flow capacity provided by the MAC-d buffer free space, that is to ensure that all active MAC-d entities are given a reasonable opportunity to use such buffer capacity to enhance their respective data flow rates. Thus, the available common transfer channel data rate is shared among all the active dedicated logical channels. The flow control is further designed to maintain the fill level of the MAC-c entity buffer at or close to a desired optimum fill level.

Referring further to Figure 5, it is seen that data transmitted by MAC-c entity 500 is directed to a transport Channel FACH.

Referring to Figure 6, there is shown flow control 530 coupled to regulate data flow into the MAC-c PDU buffers 525 from each of a number RLC's 515, each coupled to a different mobile station or other User Equipment 1-n. In a useful embodiment of the invention, the flow control operates in accordance with a "round robin" or some other principle, that is, each of the MAC-d entities connected to a given MAC-c entity are enabled to transfer data thereto sequentially, or in turn, or by flow activity (greedy manner). In this embodiment, credits are assigned to the respective MAC-d entities, in accordance with criteria described hereinafter in further detail. Upon receiving a credit, a MAC-d entity is entitled to transmit one packet data unit to the MAC-c buffer 525.

-23-

Referring to Figure 7, there is shown MAC-c buffer 525, which is available for use by dedicated channels associated with MAC-d entities. The MAC-c buffer has an optimum level $Q_{c_{opt}}$, which is related to delay and throughput requirements in buffer operation. The flow control 530 endeavors to maintain such level, using feedback information received from the MAC-d entities. The flow control also functions to keep the MAC-c buffer as low as possible, so that delays related to use of a common channel will be equally distributed among different MAC-d users. If the level of buffer 525 at a particular time is Q_c , the available buffer space $Q_{c_{diff}}$ is equal to $Q_{c_{opt}} - Q_c$. The flow control 530 operates to share this space amongst MAC-d users by assigning equal numbers of credits to them, provided they have PDU's available for transmission. The flow control calculates the credit for each MAC-d entity from both MAC-d and MAC-c buffer fill levels. More specifically, the flow control establishes credit assignment criteria which is related to buffer level $Q_{c_{opt}}$, as well to levels $Q_{c_{max}}$ and $Q_{c_{un}}$, also shown in Figure 7. These criteria are as follows:

- (1) If Q_c , MAC-c buffer fill level, is less than $Q_{c_{un}}$, an infinite or unlimited credit is given to MAC-d entities.
- (2) If the current MAC-c buffer fill level Q_c is greater than $Q_{c_{un}}$ but less than $Q_{c_{opt}}$, the difference $(Q_{c_{opt}} - Q_{c_{un}})$ is divided among active MAC-d entities.
- (3) If the current MAC-c buffer fill level Q_c is greater than $Q_{c_{opt}}$, no new credits are granted except initial credits. If the credits were previously set to 'unlimited' then they are cleared by sending of a 'zero' credits to appropriate MAC-d entities.

It is anticipated that an algorithm can readily be generated to operate flow control 530 to distribute or assign credits in accordance with the above criteria. Usefully, when a particular user becomes inactive, the algorithm will continue to distribute the remaining credits to remaining users. It is to be understood that an active user is a dedicated channel service that is connected to a common channel, and has data in its RLC-d buffer. An active user changes to a passive user when its RLC-d buffer becomes empty. The algorithm provides fair sharing of the total available band width between active users which have available band widths. The algorithm also

-24-

minimizes packet delay, which is of particular importance for users operating with small packets in situations when other users have packets of large size.

In overload situations the flow control algorithm isolates different data from each other and increases RLC buffer fill level only for dedicated channels with the highest rates. The RLC buffer fill levels of other channels with unchanged rates will not be increased. Thus, the algorithm supports working of a channel switching function that can identify and switch the overloaded channels toward dedicated transport channels. The algorithm stores the user identification where it stops with the calculations, and continues with the next user identification when the flow control algorithm is used again.

The process of transferring data from the buffers of dedicated channels to buffer 525 of MAC-c entity 510, as well as the assignment or allocation of credits to respective dedicated channels by means of a sequential or round robin procedure, is illustrated in Figure 8. More specifically, Figure 8 shows steps 600, respectively numbered 1-18, which depict the sequential assignment of credits and corresponding transfer of data about with respect to several dedicated channels, respectively incorporating radio link controls RLC-1, RLC-2, and RLC-3. Each of the steps 1-12 displays a notation indicating the step number, and also certain count information in parentheses. As indicated in Figure 8 at item 620, the count information comprises the MAC-c buffer fill level Q_c , the fill level Q_d of a corresponding one of the dedicated channel buffers, and the available credit count at the conclusion of the step. Item 620 shows the notation for step 2, by way of illustration.

Referring further to Figure 8, there is shown item 610 thereof indicating that at the beginning of the procedure shown in Figure 8, the available data space Q_{diff} in MAC-c buffer 525 is equivalent to 12 credits. As stated above, assigning a credit to a dedicated channel enables the channel to transfer a specified amount of data, such as one PDU, to the MAC-c buffer. Thus, during step 1 of Figure 8, the channel of RLC-1 receives one credit. As a result, the available space in the MAC-c buffer decreases from 12 to 11, and one PDU, corresponding to the 1 credit, can be transferred from the RLC-1 buffer to the MAC-c buffer. Accordingly, the data

-25-

contents of the RLC-1 buffer which initially equivalent to 2 credits, diminishes to a Q_d count of 1.

Referring further to Figure 8, it is seen that following step 2 the RLC-2 channel has received 1 credit, the available Q_c count has decreased to 10, and the count Q_d of buffer RLC-2 has decreased to 3 from an initial count of 4. The process is similar for step 3, associated with the RLC-3 buffer. However, the round robin procedure of the embodiment of the invention is illustrated by step 4, which shows the next-following credit again assigned to the RLC-1 channel. Upon receiving the credit, the Q_d count of the RLC-1 count goes to zero, indicating that the channel has no more PDU's available. Thus, following step 5, directed to the RLC-2 channel, the next credit is assigned to the RLC-2 channel, as indicated in step 7. This is in accordance with criterion (2), set forth above. Similarly, the Q_d count for the RLC-2 channel goes to zero following step 9, whereupon all successive credits are assigned to the RLC-3 channel.

In a second embodiment of the invention, flow control 530 again operates in accordance with a round robin or with flow activity (greedy manner) credit based algorithm. Referring to Figure 9, there is shown MAC-c buffer 525, again used by dedicated channels. The Q optimum (Q_{copt}) dashed line in figure 9 represents the optimum level in regard to unwanted buffer overflow, delay variance and good throughput. The flow control endeavors to maintain the buffer at the Q_{copt} level by means of current buffer queue fill level (Q_c) and queue maximum (Q_{cmax}) threshold values, as well as by means of information pertaining to the active users routed to the MAC-c buffer. As stated above, an active user is a dedicated channel service that is connected to a common channel and has data in its RLC-d buffer. An active user changes to a passive user when its RLC-d buffer becomes empty.

In order to construct an algorithm for the flow control of the second embodiment, a credit C is calculated as follows:

$$f = \min \left(1, \frac{(Q_{cmax} - Q_c)^2}{Q_{copt}^2} \right) \text{ Eqn. (1)}$$

-26-

$$Q_{copt} = T_h + R_{cmax} (2 * T_d) \quad \text{Eqn. (2)}$$

$$C = f \frac{Q_{copt}}{N_{mac.d}} \quad \text{Eqn. (3)}$$

5 In the above equations, T_h is processing delay/handling delay, T_d is transmission delay, and R_{cmax} is FACH transport channel maximum rate for this buffer. $N_{mac.d}$ is the number of active MAC-d entities and Q_{cmax} is a buffer high level. Default value for Q_{cmax} could be $2 * Q_{copt}$.

10 The flow control has an associated preprocessing function which can be triggered to run from an incoming PDU, a timer, or both. The preprocessor checks the credit status of users, the value of Q_c and the number of active users of buffers. The flow control algorithm is not started if the Credit is greater than 0 or other predefined value.

15 Referring to Figure 10, there is shown the general operation of the flow control for the second embodiment of invention. Initially, at 0ms., Credit (C) is equal to 2, and 3 PDU's are available at dedicated MAC-d users for transfer to MAC-c. With two credits available, two PDU's can be transferred, whereupon Credit (C) becomes 0. This starts the flow-control algorithm, which provides four credits as shown in Figure 10, which are calculated from Equation (3). The flow control is not again in process, until the number of Credits again becomes 0 or other predefined value, after 20 ms.

20 Figure 11 shows the relationship between Q_c and Credits, in accordance with Equations (1) - (3). Dashed lines indicate an adjustment which is made to limit C to a maximum value.

25 Referring to Figure 12, there is shown a modification which may usefully be employed with either of the above embodiments of the invention.

 Packets from RLC-d buffers through MAC-d are first transmitted to MAC-c appropriate buffers and then to FACH transport channels. The MAC-d packets are transferred with Data Frames that also include RLC-d buffer fill level information, as

-27-

indicated by path 700 in Figure 12. The MAC-c controls user data flow from MAC-d using Control Frames that include credits, i.e., the number of packets the user may transmit, represented in Figure 12 by path 710.

5 When a MAC-d user sends its last packet to MAC-c, the packet will be transferred with a Data Frame indicating "buffer empty", shown by path 720. Accordingly, the credits for the MAC-d user will be set to initial value by the MAC-c entity. Thus, after receiving the last packet from the user, MAC-c will send a Control Frame message with an Initial Credits parameter to the MAC-d user, shown by path 730. However, sending such message after receiving the last packet from a user
10 increases transmission delay of the next-arriving packets, and also increases the number of Control Frame messages, as shown by Figure 12.

In order to achieve a smaller delay for the next packet from a user after the buffer thereof has become temporarily empty, as well to decrease the number of control messages, an Early Initial Credit granting arrangement is used. Accordingly,
15 when the flow control algorithm shows that the calculated credit is equal to or greater than the number of packets left in a particular RLC-d buffer, the flow control algorithm in the MAC-c shall add Initial credit to the calculated credit. Thus, an initial credit will be sent to the MAC-d in advance, and there will be no need to send any new credit parameter thereto after receiving the last packet. This is illustrated by Figure 13,
20 which shows that by using the Control Frame of path 740 (Credits and Initial Credits) the additional Control Frame of path 730 in Figure 12 is eliminated.

Although preferred embodiment(s) of the methods, systems, and arrangements of the present invention have been illustrated in the accompanying Drawings and described in the foregoing Detailed Description, it will be understood that the present
25 invention is not limited to the embodiment(s) disclosed, but is capable of numerous rearrangements, modifications, and substitutions without departing from the spirit and scope of the present invention as set forth and defined by the following claims.

-28-

WHAT IS CLAIMED IS:

1. A method in a wireless communications system for providing flow control, comprising the steps of:

- 5 providing a receiver entity;
- providing at least one transmitter entity, said at least one transmitter entity transmitting when said at least one transmitter entity has credit from said receiver entity;
- determining whether to give credit to said at least one transmitter entity;
- 10 and
- giving credit to said at least one transmitter entity by said receiver entity if it is determined in said step of determining to give credit to said at least one transmitter entity.

- 15 2. The method according to Claim 1, further comprising the step of:
- transmitting a packet from said at least one transmitter entity to said receiver entity when said at least one transmitter entity has credit, transmitting of said packet being carried out in a round robin fashion, or in accordance with a flow activity (greedy manner) credit based principle.

- 20 3. The method according to Claim 1, wherein said receiver entity comprises a Medium Access Control (MAC)-c entity, and said at least one transmitter entity comprises a MAC-d entity, and said wireless communications system comprises a Universal Mobile Telecommunications System (UMTS).

- 25 4. The method according to Claim 1, wherein said receiver entity comprises a receiver buffer, and said at least one transmitter entity is associated with another buffer.

-29-

5. The method according to Claim 4, wherein the determination in said step of determining is based, at least in part, on both a fill level of said receiver buffer and a fill level of said another buffer.

5 6. The method according to Claim 4, wherein the determination in said step of determining is based, at least in part, on a fill level of said receiver buffer, but not on a fill level of said another buffer.

7. A method in a wireless communications system for providing flow control, comprising the steps of:

10 providing a receiver entity, said receiver entity having a receiver buffer with a fill level;

15 providing a plurality of transmitter entities, each one of said plurality of transmitter entities transmitting when said each one of said plurality of transmitter entities has been given credit from said receiver entity;

analyzing said fill level of said receiver buffer with respect to a first level, a second level, and a third level; and

giving credit from said receiver entity responsive to said step of analyzing.

20

8. The method according to Claim 7, wherein said step of analyzing comprises the step of determining whether said fill level is less than said first level, and if so, said step of giving credit comprises the step of giving infinite credit to all of said plurality of transmitter entities.

25

9. The method according to Claim 7, wherein said step of analyzing comprises the step of determining whether said fill level is greater than said first level but less than said second level, and if so, said step of giving credit comprises the step of giving the difference between said third level and a current level in credit divided over all active transmitter entities of said plurality of transmitter entities.

30

-30-

10. The method according to Claim 9, wherein said difference between said third level and said current level comprises an available credit amount and said step of giving the difference between said third level and a current level comprises the steps of:

5 ascertaining if a first active transmitter entity of said all active transmitter entities of said plurality of transmitter entities has a current credit of less than a buffer fill level of an associated buffer;

 if so, increasing said current credit of said first active transmitter entity by one and decreasing said available credit amount by one;

10 repeating said step of ascertaining for each active transmitter entity of said all active transmitter entities of said plurality of transmitter entities until either said available credit amount is exhausted or said all active transmitter entities of said plurality of transmitter entities has their corresponding current credit equal to their corresponding buffer fill level; and

15 recording a last active transmitter entity of said all active transmitter entities of said plurality of transmitter entities to receive a credit and making the active transmitter entity after said last active transmitter entity the next first active transmitter entity.

20 11. The method according to Claim 7, wherein said step of analyzing comprises the step of determining whether said fill level is greater than said second level, and if so, said step of giving credit comprises the step of giving no new credits to any of said transmitter entity except initial credits.

25 12. The method according to Claim 11, wherein if the credits were previously set to 'unlimited', they are cleared by sending a zero credits message to corresponding transmitter entities.

30 13. A method in a wireless communications system for providing flow control, comprising the steps of:

-31-

providing a receiver entity, said receiver entity having a receiver buffer with a fill level;

providing a plurality of transmitter entities, each one of said plurality of transmitter entities transmitting when said each one of said plurality of transmitter entities has been given credit from said receiver entity;

analyzing said fill level of said receiver buffer with respect to a first level, a second level, and a number of active transmitter entities of said plurality of transmitter entities; and

giving credit from said receiver entity responsive to said step of analyzing.

14. The method according to Claim 13, wherein said step of analyzing comprises the step of determining whether said fill level is greater than said second level, and if so, said step of giving credit comprises the step of giving no new credit to all active transmitter entities of said plurality of transmitter entities, but said all active transmitter entities of said plurality of transmitter entities are permitted to transmit outstanding credit.

15. The method according to Claim 13, wherein said step of analyzing comprises the step of determining whether said fill level is less than said second level, and if so, said step of giving credit comprises the step of giving credit to all active transmitter entities of said plurality of transmitter entities in accordance with a value equal to a function times a quotient of said first level divided by said number of active transmitter entities of said plurality of transmitter entities, said function equal to a minimum of one and a quotient of an n^{th} -power of a difference between said second level and said fill level divided by an n^{th} -power of said first level.

16. The method according to Claim 15, wherein the variable n of said n^{th} -power is equal to 2.

30

-32-

17. A method in a wireless communications system for providing flow control, comprising the steps of:

receiving a data frame, said data frame including a buffer fill level indicator; and

5 transmitting a control frame responsive to said data frame, said control frame including a credits provision, said credits provision equal to a sum of an ordinary credits provision and an initial credits provision.

18. The method according to Claim 17, further comprising the steps of:

10 receiving a second data frame, said second data frame including a buffer empty indicator; and

receiving a third data frame, said third data frame including a buffer fill level indicator.

15 19. The method according to Claim 17, wherein said step of transmitting a control frame is performed responsive to a determination that said credits provision is equal to or greater than said buffer fill level.

20 20. A method in a wireless communications system for providing flow control, comprising the steps of:

providing a receiver entity;

providing at least one transmitter entity, said at least one transmitter entity having a transmitter buffer with a buffer fill level; and

25 comparing a credits provision that is allocated to said at least one transmitter entity with said buffer fill level.

21. The method according to Claim 20, further comprising the steps of:

determining if said credits provision is equal to or greater than said buffer fill level responsive to said step of comparing; and

-33-

if so, transmitting from said receiver entity to said at least one transmitter entity a control frame, said control frame including said credits provision and an initial credits provision.

5

1/8




FIG. 1

200

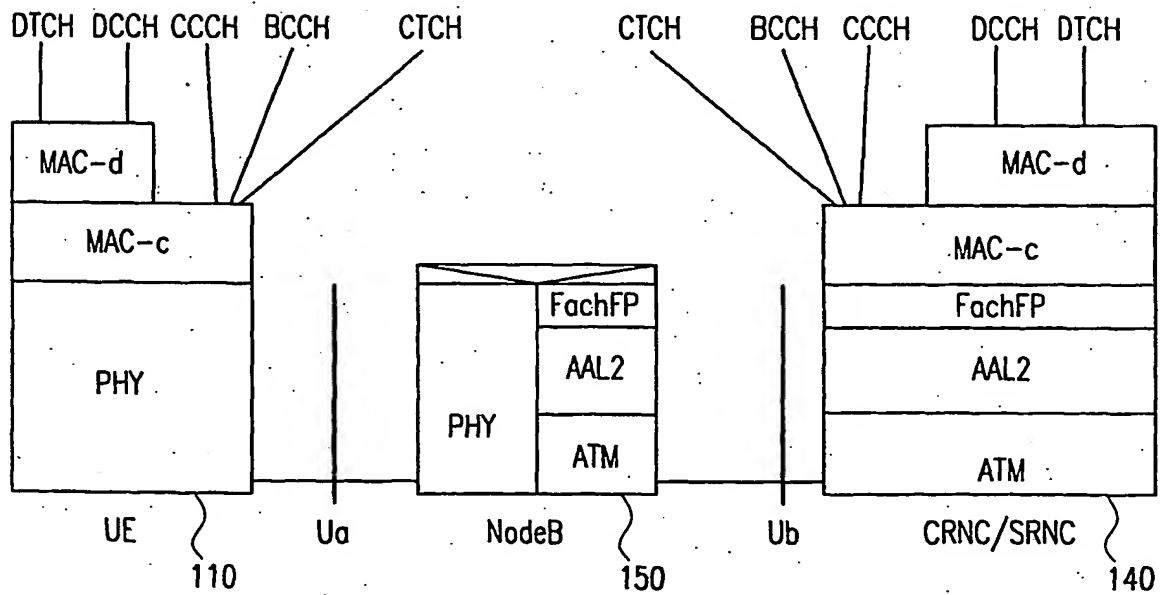


FIG. 2


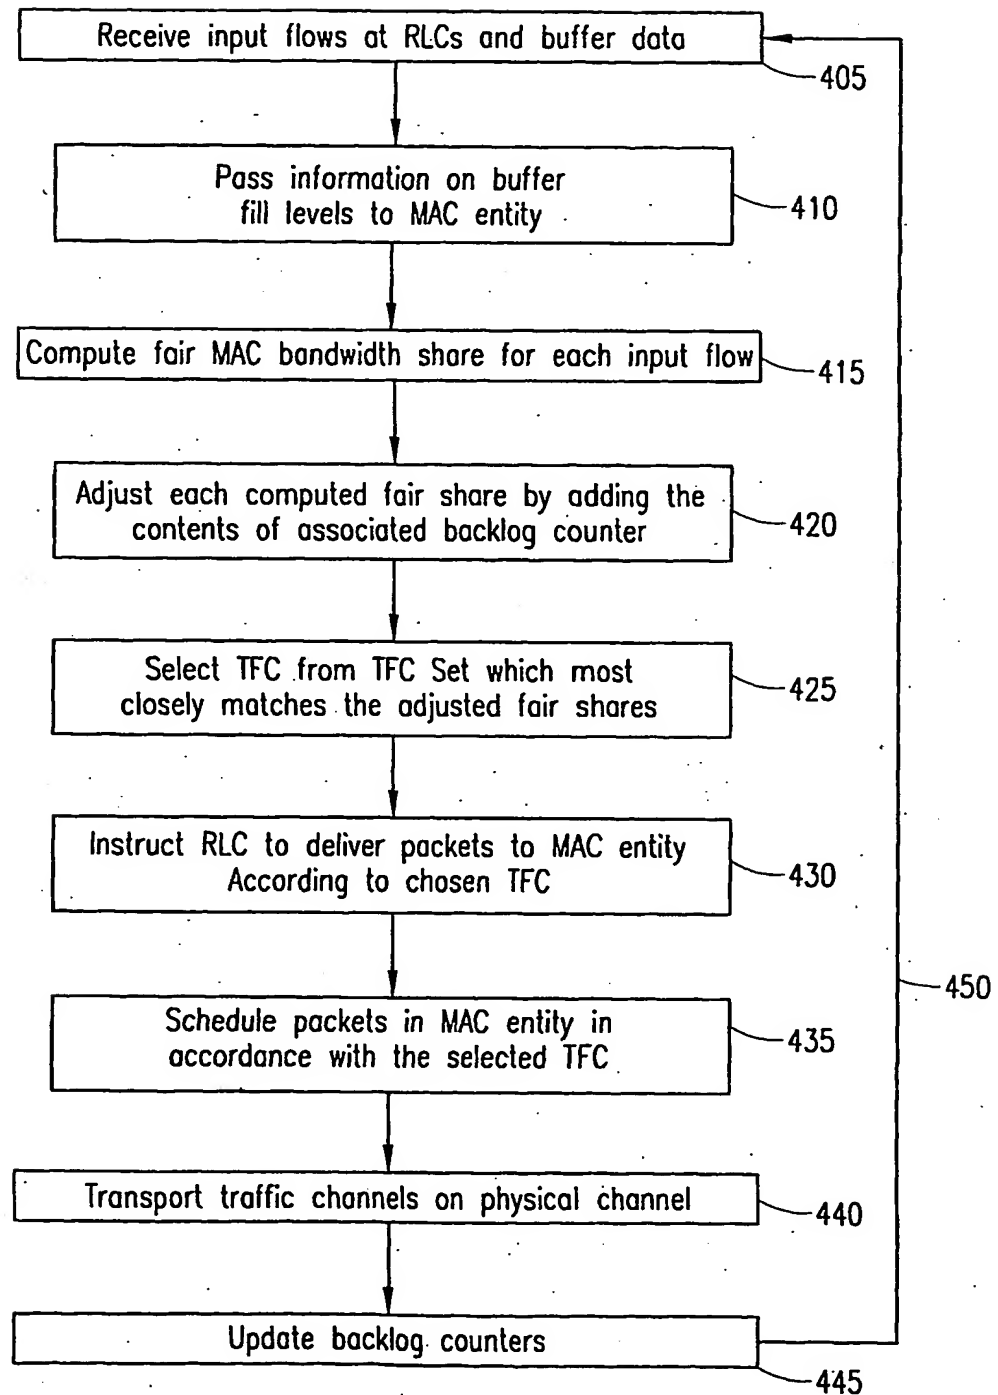


FIG. 3

3/8

400*FIG. 4*




FIG. 5




FIG. 8

5/8




FIG. 6

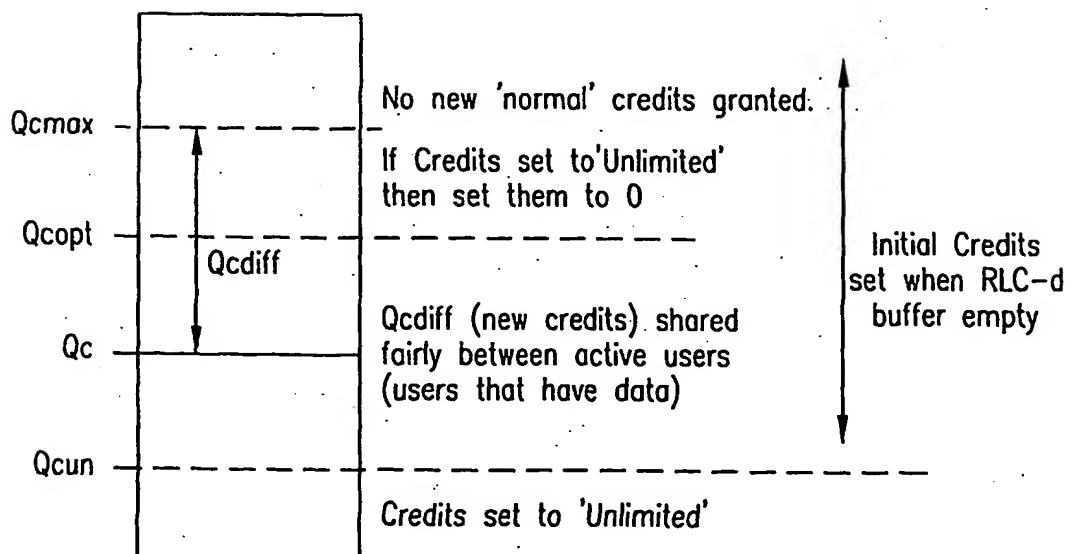




FIG. 7

6/8

*FIG. 9**FIG. 11*

7/8

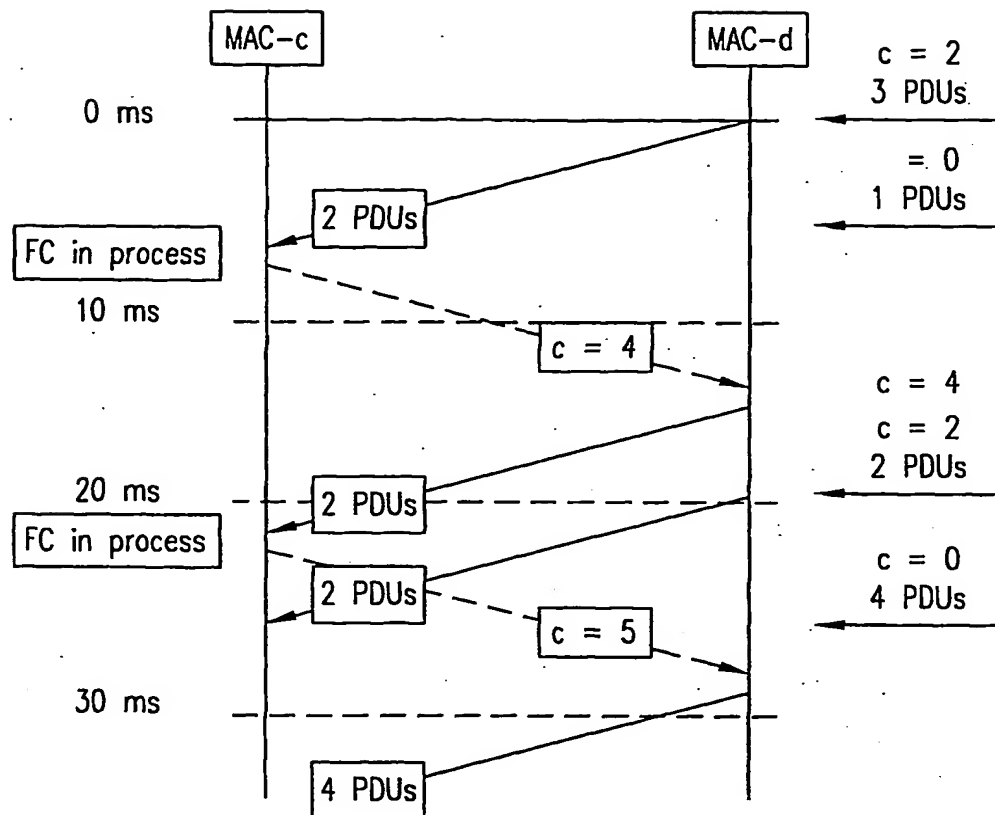


FIG. 10

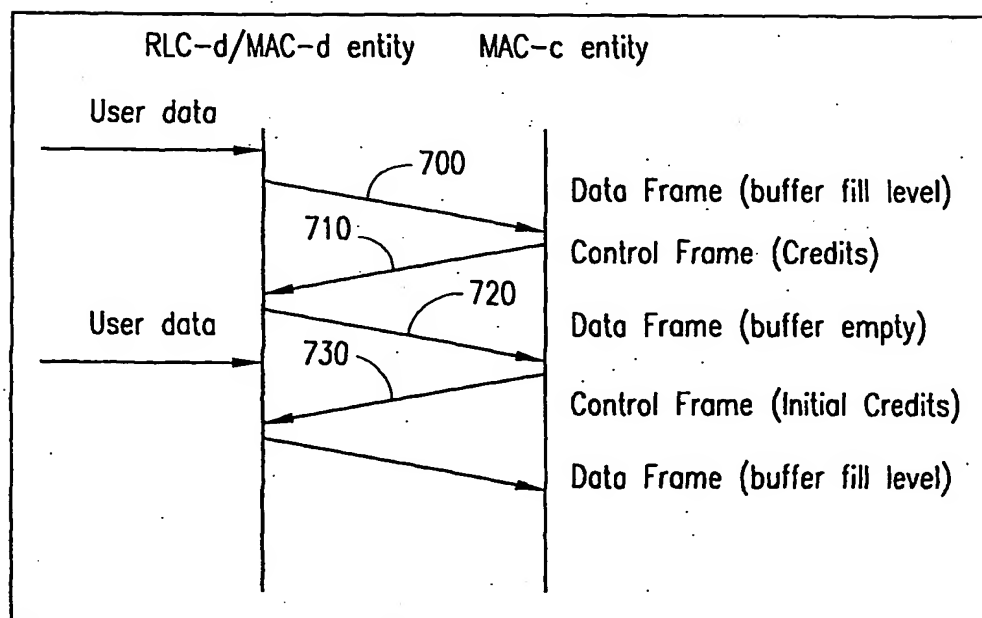



FIG. 12

**FIG. 13**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 01/00407

A. CLASSIFICATION OF SUBJECT MATTER

IPC7: H04L 12/56, H04Q 7/22

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC7: H04L, H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5748613 A (KILK ET AL), 5 May 1998 (05.05.98), column 1, line 54 - column 2, line 49, claims 1-18, abstract	1,2,4-21
Y	--	3
X	US 5831985 A (SANDORFI), 3 November 1998 (03.11.98), column 1, line 56 - line 61; column 2, line 6 - line 52, figure 4	1,2,4-21
Y	--	3



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

27 June 2001

Date of mailing of the international search report

13. 07. 2001

Name and mailing address of the International Searching Authority
European Patent Office P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk
Tel(+31-70)340-2040, Tx 31 651 epo nl,
Fax(+31-70)340-3018

Authorized officer

Elisabet Aselius/JAN
Telephone No.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 01/00407 ..

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 9700568 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION), 3 January 1997 (03.01.97), page 2, line 4 - line 10, figures 2a, 2b, abstract	1,2,4-21
Y	--	3
P,X	WO 0042792 A1 (NOKIA NETWORKS OY), 20 July 2000 (20.07.00), page 4, line 1 - line 16; page 4, line 30 - page 5, line 10; page 6, line 8 - line 36, page 7 line 7-20, page 9 line 6-16, page 9 line 26 - page 10 line 21, page 12 line 1-17, abstract	1-7
P,A	--	8-21
P,X	WO 0041365 A1 (MICROSOFT CORPORATION), 13 July 2000 (13.07.00), figures 4-8, abstract	1,2,4-21
P,A	US 6108713 A (SAMBAMURTHY ET AL), 22 August 2000 (22.08.00), column 3, line 32 - line 64; column 4, line 20 - line 31; column 4, line 44 - line 53, col 4 line 65- col 5 line 15, col 7 line 53-col 8 line 6, col 10 line 41-54, col 12 1 57-col 36 line 19, fig 2,3,10, abstract	1-21
	-- -----	

SA 327987

INTERNATIONAL SEARCH REPORT

Information on patent family members

28/05/01

International application No.

PCT/SE 01/00407

Patent document cited in search report			Publication date	Patent family member(s)		Publication date
US	5748613	A	05/05/98	DE	19652541 A	02/10/97
				GB	2311700 A,B	01/10/97
				GB	9705174 D	00/00/00
				JP	10023078 A	23/01/98

US	5831985	A	03/11/98	NONE		

WO	9700568	A1	03/01/97	DE	69509804 D,T	09/12/99
				EP	0832527 A,B	01/04/98
				ES	2132669 T	16/08/99
				JP	10511829 T	10/11/98
				KR	257437 B	15/05/00
				PL	178412 B	28/04/00
				PL	324002 A	27/04/98
				US	6181683 B	30/01/01

A1	0042792	A1	20/07/00	NONE		

WO	0041365	A1	13/07/00	AU	3128100 A	24/07/00

US	6108713	A	22/08/00	AU	6656998 A	26/08/98
				WO	9835307 A	13/08/98
				AU	6169398 A	26/08/98
				EP	0960517 A	01/12/99
				US	6076115 A	13/06/00
				US	6085248 A	04/07/00
				WO	9835480 A	13/08/98

THIS PAGE BLANK (USPTO)